

Mars Exploration Rovers Orbit Determination Filter Strategy

Timothy P. McElrath^{*}, Michael M. Watkins[†], Brian M. Portock[‡], Eric J. Graat[§], Darren T. Baird^{**},
Geoffrey G. Wawrzyniak[§], Joseph R. Guinn^{††}, Peter G. Antreasian^{‡‡}, Amy A. Attiyah^{§§}, Ronald C. Baalke^{§§},
William L. Taber^{***}

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, 91109

The successful delivery of the Mars Exploration Rover (MER) landers to well within the boundaries of their surface target areas in January of 2004 was the culmination of years of orbit determination analysis. The process began with a careful consideration of the filter parameters used for pre-launch covariance studies, and continued with the refinement of the filter after launch based on operational experience. At the same time, tools were developed to run a plethora of variations around the nominal filter and analyze the results in ways that had never been previously attempted for an interplanetary mission. In addition to achieving sub-kilometer Mars-relative orbit determination knowledge, the filter strategy and process detected unexpected error sources, while at the same time proving robust by indicating the correct solution. Consequently, MER orbit determination set a new standard for interplanetary navigation.

Nomenclature

ΔV	=	delta velocity
ΔDOR	=	delta-Differential One-way Range, an angular measurement made using VLBI techniques
$VLBI$	=	Very Long Baseline Interferometry
$nrad$	=	nanoradians
$nsec$	=	nanosecond, about 30 cm at the speed of light
$B \bullet R$	=	component of B-vector (hyperbolic miss vector) along R axis
$B \bullet T$	=	component of B-vector (hyperbolic miss vector) along T axis

I. Introduction

The Mars Exploration Rover project launched two nearly identical rovers to Mars in the summer of 2003. Due to the tight Mars delivery requirements necessary for a successful landing, the orbit determination of these spacecraft challenged the capabilities of the navigation team. Whereas a more complete description of the mission can be found in Ref. 1, the main aspects of the mission relevant to the filter strategy are that the spacecraft were spin-stabilized, and that the tracking data available consisted of Doppler, range, and ΔDOR measurements from the NASA-JPL Deep Space Network (DSN) tracking stations. The orbit determination results, which are summarized in Ref. 1 and 2, are most commonly expressed in terms of the Mars B-plane, but the most significant parameter is the entry flight path angle (EFPA), which is closely related to the B vector magnitude. This paper consists of four main sections: baseline filter selection, filter variations, statistical tests and presentation, and results.

^{*} Inner Planet Navigation Group Supervisor, Navigation and Mission Design Section, JPL, M/S 301-276

[†] Mission System Manager, Mars Science Laboratory Project, JPL, M/S T1723-118

[‡] Member of Engineering Staff, Navigation and Mission Design Section, JPL, M/S 301-276, Member AIAA

[§] Member of Engineering Staff, Navigation and Mission Design Section, JPL, M/S 264-380

^{**} Member of Engineering Staff, Navigation and Mission Design Section, JPL, M/S 264-380, Member AIAA

^{††} Mission Manager, Phoenix Project, JPL, M/S 264-380, Member AIAA

^{‡‡} Member of Engineering Staff, Navigation and Mission Design Section, JPL, M/S 230-205

^{§§} Member of Engineering Staff, Navigation and Mission Design Section, JPL, M/S 301-125L

^{***} Navigation Software Group Supervisor, Navigation and Mission Design Section, JPL, M/S 301-125L

II. Baseline Filter Selection

The orbit determination process requires the development of correct statistical models for the spacecraft dynamics and the tracking measurement errors. Due to the relative simplicity of the MER spacecraft, the modeling of the Earth-based measurements comprises the majority of the filter elements, and is also the most useful part for other spacecraft, for which the tracking system is essentially identical. The measurement model consists of white data noise (also referred to as the data weight), and of measurement parameters for which partial derivatives and statistical models exist. The latter are used to model long-period effects that are significantly different from white noise. In an inadvertent and unwarranted homage to Ptolemy, the description of the measurement modeling parameters will proceed from the center of the Earth to Mars, before concluding the section with a discussion of data weighting and of spacecraft trajectory modeling.

For all of the parameters (*i.e.* everything but data weights), the filter allows each to be designated as a constant estimated parameter (with *a priori* value and covariance), a stochastic parameter (with an update time sequence, each item of which includes a correlation time (if any) and a process noise), or a consider parameter (with a covariance that is allowed to affect the estimate without being estimated itself). The same parameter can be estimated both as a constant and as a stochastic, which has the effect of removing a bias from the stochastic estimates. Consider parameters model the effect of a constant offset in a parameter over all the data, which makes their use for time-varying errors pessimistic, despite being convenient. Conversely, treating a parameter as a stochastic necessarily allows it to change, which may cause it to compensate for errors in other models in an unrealistic way. As a consequence, stochastic estimates must be monitored, even if the expected behavior is not interesting in itself.

Before launch, an examination of all of the modeling factors covered in this paper produced a baseline set of filter assumptions, which in turn was used to produce the pre-launch delivery covariance predictions. These pre-launch assumptions are shown in Table 1 (reproduced from Ref. 1). During flight, a number of changes were made to the assumptions, both as a result of in-flight experience and due to continued research into error models. The final approach assumptions are shown in Table 2 (also reproduced from Ref. 1). Both tables will be referenced throughout this section to illustrate the continuing development of the filter.

A. Earth Orientation

The Doppler and range data determine the position of the spacecraft with respect to the assumed locations of the tracking stations, which in turn depend on the orientation of the Earth. Whereas the long-period effects of precession, nutation, and tidal oscillations are well determined, the polar motion and rotational phase (expressed by UT1) are essentially random after only a few days. Consequently, there are a variety of measurements made to maintain the knowledge of these parameters, and improve their prediction. However, due to latencies in the data, and the stochastic nature of the process being observed, the steady-state uncertainties are only reached for points several days in the past, and are significantly worse at the end of the data arc.

The Earth orientation model values used by JPL navigation teams are produced by the Tracking Systems and Applications Section at JPL, based on inputs from JPL and non-JPL sources. The error characteristics and details of the input data are described in Ref. 3. Due to unavoidable processing latencies, the Earth orientation model file became available no earlier than about 11am PST, with data through 4pm PST the previous day, resulting in almost 1 day of latency at best. However, the most recent data (which is all GPS) is not as useful as the VLBI intensive series from Goddard Spaceflight Center (GSFC), which takes several days to process. The normal evolution of Earth orientation errors is shown in Fig. 1a.

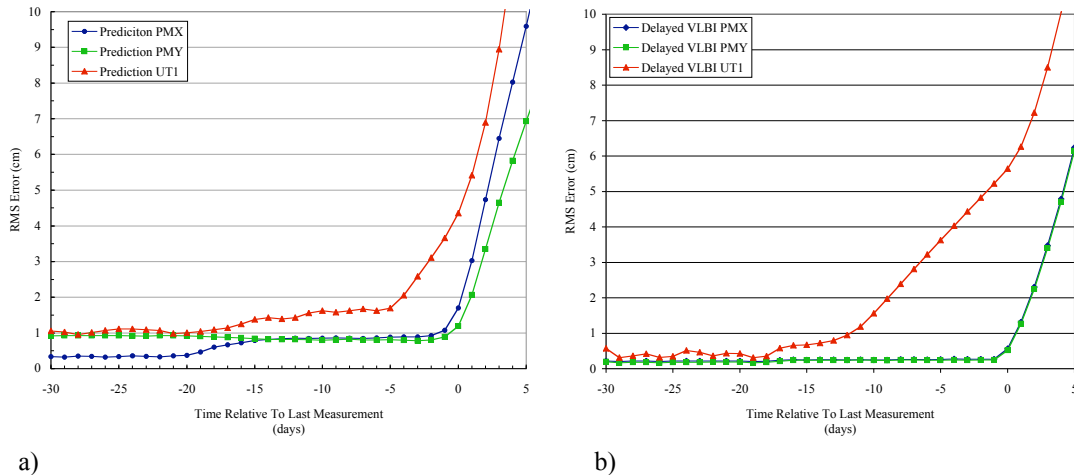
Before launch, the Earth orientation model began with a process noise of 2 centimeters, with 4 stochastic updates per day until 7 days before the end of the arc. At that point, the process noise was increased to 10 centimeters by the end of the data arc in linear steps once per day. Despite being an improvement on previous models, the values used were general numbers based on experience rather than resulting from a specific study. Once Ref. 3 was available, the model was updated with a 2-day correlation time, and increased process noise values at epochs chosen such that the resulting sigmas matched the predicted sigmas. Since the knowledge of UT1 degraded more quickly than that of the polar motion, the increase in UT1 uncertainty started earlier. The resulting error model sigmas are shown in Fig. 2.

Table 1: OD filter configuration – pre-launch.

Error Source	Estimated?	A Priori Uncertainty (1 σ)	Correlation Time	Update Time	Comments
2-way Doppler (mm/s)	–	0.075	–	–	~4.5 mHz
Range (m)	–	4	–	–	29 range units
Δ DOR (nrad)	–	4.5	–	–	0.12 ns
Epoch State					
Position (km)	✓	1000	–	–	
Velocity (km/s)	✓	1	–	–	
Solar Pressure					
Area (%)	✓	5	–	–	Sunlit area of spacecraft.
Specular & Diffuse Coefficients (%)	✓	10	–	–	
ACS Event ΔV (mm/s)					
Line-of-Sight Comp.	✓	3	–	–	Every 8 days.
Lateral Comp.	✓	3	–	–	
Normal Comp.	✓	3	–	–	
TCMs					
TCM-1	✓	422, 440	–	–	MER-A, MER-B respectively
TCM-2	✓	17, 15	–	–	Spherical uncertainty (mm/s)
TCM-3	✓	3, 5	–	–	TCM-4 at E - 8 days
TCM-4	✓	3, 3	–	–	TCM-5 at E - 2 days
TCM-5	✓	3, 3	–	–	TCM-6 at E - 6 hours
TCM-6	✓	7, 7	–	–	5% (3 σ) proportional error (per axis)
Non-gravitational Accelerations (km/s ²)	✓	2.0×10^{-12}	10 days	1 day	6 mm/s (3 σ) fixed error (per axis)
Range Bias (m)	✓	2	0	Per pass	Spherical covariance; estimated daily (1 day batches).
Doppler Bias (mm/s)	✓	0.005	0	Per pass	Estimated per pass.
Mars & Earth Ephemerides	–	DE405+	–	–	Estimated per pass.
Station Locations (cm)	–	3	–	–	
Quasar Locations (nrad)	–	2	–	–	
Pole X, Y (cm)	✓	2 – 10	0	6 hours	*Use lower value up to 7 days before data cutoff; ramp up to higher value at data cutoff.
UT1 (cm)	✓	2 – 10	0	6 hours	(UT1: 0.256 ms = ~10 cm.)
Ionosphere – day (cm)	✓	55	0	6 hours	S-band values.
Ionosphere – night (cm)	✓	15	0	6 hours	
Troposphere – wet (cm)	✓	1	0	6 hours	
Troposphere – dry (cm)	✓	1	0	6 hours	

Table 2: OD filter configuration – final approach.

Error Source	Estimated?	A Priori Uncertainty (1σ)	Correlation Time	Update Time	Comments
2-way Doppler (mm/s)	–	Weight by pass: ≥0.05	–	–	3.36 x RMS (60 sec) of pass-through residuals. (1.0 mm/s = ~0.056 Hz.)
Range (m)	–	Weight by pass: ≥0.14	–	–	2.22 x RMS of pass-through residuals. (1.0 m = ~7.0 RU.)
ΔDOR (nrad)	–	~3.2	–	–	0.085 ns
Epoch State					
Position (km)	✓	100,000	–	–	Effectively infinite.
Velocity (km/s)	✓	1.0	–	–	Effectively infinite.
Solar Pressure					
Specular Coefficients	✓	0.05	–	–	All components; 0.1 in normal units.
Diffuse Coefficients	✓	0.033			All components; 0.1 in normal units.
Solar Array Diffuse Coef.	✓	0.033	7 days	1 day	Correlation broken at turns.
ACS Event ΔV (mm/s)					
Line-of-Sight Comp.	✓	0.05, 0.1	–	–	MER-A & MER-B values respectively, based on results of ACS/NAV calibration. Location per current plan; all future events included.
Lateral Comp.	✓	0.05, 0.05	–	–	
Normal Comp.	✓	0.05, 0.05	–	–	
TCM-4					
Magnitude (%)	✓	1.67	–	–	Equivalent to 0.127 N.
Pointing (deg)	✓	0.5	–	–	Use 0.5 deg for both cone and clock.
Non-gravitational Accelerations (km/s ²)	–	–	–	–	No longer used – accounted for in solar pressure uncertainties and modeling of ACS turns as discreet events.
Charged-particle Effects					
Delay (m)	✓	1.5	–	1 day	
Delay Rate (m/s)	✓	0.00005	–	1 day	Equivalent to ~4 meters per day
Range Bias (m)	✓	2	0	Per pass	Estimated per pass.
Doppler Bias (mm/s)	–	–	–	–	No longer used.
Mars & Earth Ephemerides	–	DE410	–	–	0.5 x DE405+; Mars RTN 1σ uncertainties: 9 m, 136 m, 442 m.
Station Locations (cm)	–	Per covariance.	–	–	Use latest updates to station locations and covariance.
Quasar Locations (nrad)	–	2	–	–	
Pole X, Y (cm)	✓	1, 4	2 days	6 hours	Use larger value for last 2 – 6 days per KEOF memo. (For UT1, 0.256 ms ⇒ ~10 cm.)
UT1 (cm)	✓	1.7, 9	2 days	6 hours	
Ionosphere – day (cm)	✓	55, 300	6 hours	1 hour	Subsequent passes uncorrelated; use larger values (6X) for ionosphere when only predicted calibrations available.
Ionosphere – night (cm)	✓	15, 100	6 hours	1 hour	
Troposphere – wet (cm)	✓	1, 1	6 hours	1 hour	
Troposphere – dry (cm)	✓	1, 1	6 hours	1 hour	S-band units.



a) **Figure 1. EOP error comparison between normal data latency and delayed VLBI processing.**
 a) EOP error assuming normal data latency.
 b) EOP error assuming delayed VLBI processing.

The last week before the landing of MER-A coincided with the holiday season at the end of 2003. The VLBI intensive series processing at GSFC was delayed since there was no contractual agreement to provide the data with the usual latency, although efforts were made to reduce the processing backlog once the importance of the data was known. Once it was understood that the Earth orientation (and particularly UT1) knowledge was degraded, an alternate set of process noise values was developed, with an assumed latency of an additional day. The formal sigmas from the Earth orientation estimation process with delayed processing are also shown in Fig. 1b. Once past New Years Day 2004, the Earth orientation latency did not become an issue again. As mentioned above, treating Earth orientation stochastically required the periodic inspection of the resulting estimates, which were able to detect UT1 errors when Δ DOR data was included.




Figure 2. EOP *a priori* sigma model based on Ref. 3 data

B. Station Locations

While Doppler and range data are sensitive to station locations (and through them to Earth orientation), Δ DOR data are much less sensitive, due to their use of an external reference (the quasar), which corrects almost all Earth-based modeling errors, especially as the quasar-spacecraft separation decreases to zero, as described in Ref. 2. However, accurate modeling of the station locations is still important for deriving an accurate solution with the Δ DOR data withheld, which is important in evaluating solution consistency. The DSN stations used by MER are located in three tracking complexes near Goldstone, CA; Madrid, Spain; and Canberra, Australia, and comprise 5, 4, and 3 antennas at each site, respectively. The station locations are developed from a combination of survey data within the DSN complex and VLBI data, primarily between complexes, but also within them. The result of the station location solution is a set of coordinates and a correlated covariance, which is scaled to account for systematic effects. Typical sigmas are 3 to 4 centimeters in each direction, with correlations of about 0.6 between complexes and close to 1 within complexes. Two exceptions were for DSS 26 at Goldstone, which began tracking operations shortly before launch, and DSS 55 at Madrid, which began tracking during the middle of cruise. Both of these stations consequently had significantly less VLBI data available to determine their location. A pre-launch DSS 26 location was found to be in error by almost a meter vertically, and so a new station location set incorporating the

corrected DSS 26 location and the newly-determined DSS 55 location was developed during cruise. Nonetheless, both stations had locations that were known to within 5 centimeters, and the difference did not contribute significantly to the resulting delivery covariance. Since the station location error was very nearly constant over the timescale of the MER cruise, and since the tracking data (even including Δ DOR data, since only 1 or 2 quasars were used in each observation session) could not improve the station location knowledge, the station location errors were one of the few consider parameters.

C. Range Biases

Since a range measurement must be made through all of the tracking station uplink and downlink electronics, additional delays due to path length and electronic frequency response are introduced. During the pre-pass calibration, the tracking station contribution to the total delay is measured, but the result is only good to perhaps half a meter of random error at best, with an additional constant delay of about 1 meter, due to factors such as the inability of the calibration to introduce a realistic Doppler shift on the ranging signal. Pre-launch modeling used 2 meters as a per-pass bias (with no correlation from pass to pass), and no constant bias, based on then-current conservative accuracy assumptions. Since the primary effect of the range bias assumptions is on solutions that do not contain Δ DOR data, accurate range bias modeling is mainly desired to produce solution consistency, and conservatism doesn't degrade the delivery results.

Although the middle part of cruise had fairly sparse tracking (3-4 passes/week), the final 45 days before each arrival had 2-3 tracks per day, or almost continuous tracking. Since enough Δ DOR data were available to precisely determine the geocentric angular position of the spacecraft, the dense tracking coverage clearly showed small trends in the range biases. With the luxury of having two spacecraft being simultaneously (over a few-day timescale) tracked by the same combination of stations, comparisons were made of the per-station biases between the spacecraft. Whereas the results were not identical (which is not surprising given the 4 MHz difference in downlink frequencies and the presence of frequency-dependent delays), several stations showed remarkable similarity in constant range biases. Accordingly, range biases of up to 1.5 meters were applied from one spacecraft's estimate to the other spacecraft (to avoid inconsistent reuse of the same data) for DSS 26, 43, 54, and 55^{†††}, and a solution that estimated constant station range biases (always in addition to the per-pass biases) was added to the battery of solutions, from which it began to be chosen as the delivered solution on a regular basis.

In addition to the small biases discussed above, occasionally passes would be biased by a larger amount (3 to 10+ meters), primarily due to difficulties executing the pre-pass range calibration. Such passes would have their per-pass bias *a priori* sigma increased to 10 meters or more, and possibly have their values corrected to be consistent to nearby passes. Regardless of the adjustment, the effect of the large sigmas was to keep the bias for that pass from adversely affecting the solution. The spacecraft also contributes a delay to the ranging signal, primarily due to transponder electronics. This delay is calibrated before launch and assumed to be constant, and in practice does not need to be estimated due to other errors with similar appearance (*e.g.*, the Mars ephemeris error) being present in the filter.

D. Transmission Media

The radio signal propagating between the tracking station and the spacecraft experiences delays due to the Earth's troposphere and ionosphere, and also due to interplanetary solar plasma. The troposphere is modeled using meteorological data at the tracking complexes, combined with GPS-derived estimates of the troposphere difference in the upper atmosphere. The result of the troposphere measurements is a model for zenith delay, which is then mapped to the elevation angle of each tracking measurement using the Niell mapping model. The ionosphere is calibrated by processing all available GPS data to form a global model, and then evaluating the model along the ray path to each spacecraft. The interplanetary plasma can most effectively be measured with a dual-frequency signal (which MER lacked) for each spacecraft, but differences in its effect on range and Doppler data can be exploited to remove some of the effect.

The available troposphere parameters include the wet and dry components for each station complex, and the ionosphere parameters include day and night components for each complex.^{†††} The remaining uncalibrated error is expressed as a combination of the two components for each media type, as shown in Table 1 and 2. Since these errors change for each pass, they must be estimated as stochastics, even though the parameter partials are only

^{†††} The most likely explanation for these biases are errors in determining the delay between the end of the calibration path and the physical station location. Note that the two newest stations showed up in this list again.

^{††††} Ionosphere delay is expressed as zenith delay at S-band (2295 MHz), which produces a delay 11.65 times larger than the average of the MER uplink and downlink X-band frequencies using a frequency-squared dependency.

approximate. Pre-launch treatment comprised a 6-hour update time with no correlation, but this is not an especially realistic model. The measurement sensitivity to Earth media increases markedly at lower elevation angles (as is reflected in the partials), so the most important part of the model is to allow for differences between the rising and setting part of the pass. The adopted model used 1-hour updates with a 6-hour correlation time for each of the four Earth media parameters for each pass, with no correlation between passes. Although this model is still somewhat *ad hoc*, it is at least approximately realistic, and worked reasonably well for MER. Nonetheless, the Earth media estimation model would benefit from more study than the MER Navigation team was able to devote to it.

Although the range and Doppler data are affected more severely by Earth media, the Δ DOR data is not completely immune to temporal (on the timescale of scan separation) or spatial (depending on the quasar-spacecraft angular separation) errors, even after all of the differencing has been performed. The media partials are correctly computed for the Δ DOR data as well, and the parameters were treated as constants over the entire Δ DOR observation (almost always consisting of two points separated by less than an hour). Care was taken to insure that the stochastic media estimate from the Doppler and range data was not correlated with the Δ DOR media estimates, since otherwise any problems with the Doppler and range data (whether related to transmission media or not) could corrupt the Δ DOR data.

While the Earth orientation parameters suffer a gradual increase in uncertainty after the last data, the effect on the media calibrations is more of a step function. Consequently, after the last calibrated pass the Doppler and range tracking data were downweighted to 3 times the nominal weights (or 0.225 mm/sec and 3 meters, respectively), and the stochastic process noise was increased, to a factor of 6 for the ionosphere and a factor of 2 for the troposphere.

The middle of the MER cruise period was marked by unusually large solar flares and related solar activity, to the point of affecting spacecraft attitude sensors and causing concern about permanent damage to the electronics. The navigation effect was most noticeably to introduce noise bursts into the Doppler data². However, in December 2003, range biases increasing to several meters were observed, most notably on December 18-19. This is shown in Fig. 3, which displays range residuals without the estimated range biases having been applied. The particular event shown here persisted for about two days (or about a day beyond what is plotted in Fig. 3).

Whereas a nearly constant range bias due to steady-state solar plasma abundance is not a concern (similar to the transponder bias discussed above), changes in solar plasma produce a range (group) delay but carrier-phase advance (causing Doppler shifts opposite to the range data rate), which is impossible to fit by any kinematic trajectory adjustment. However, the uniqueness of the signature makes it possible to estimate the charged-particle error without much risk of other effects becoming aliased into the new parameters.

Once the solar plasma effect was identified, two models were introduced to account for it. The first was a simple model using a random-walk range bias and a white-noise Doppler bias, with a process noise large enough to respond to the observed biases. The second model (referred to hereafter as the ‘charged-particle model’) used Hermite interpolation on one-day centers to obtain a continuous range and Doppler effect with the appropriate sign difference. The estimated parameters were the range delay and range delay rate at each epoch (chosen to be once per day at 12:00 UTC). The initial *a priori* sigmas were 4 meters and 4 meters/day, but the range value was later reduced to 1.5 meters, as will be described below. Both of these models worked reasonably well, although the random-walk approach somewhat unreasonably increased the geocentric range errors to tens of meters. The charged-particle model, which is described in more detail in Ref. 4, became the standard filter by the last week of the MER-A approach, due to the improved consistency of the Doppler and range-only solutions incorporating it.

One difficulty that arose with the charged-particle model came to light during the period of increased UT1 errors in late December, 2003. Before the UT1 errors were known, the range delay estimate from the charged-particle model started increasing linearly in an alarming manner, to a value of several meters. Earth-based measurements of




Figure 3. Range data residuals showing solar plasma growth.

the solar plasma density, which had earlier shown good consistency with the MER charged-particle estimates⁴, did not record such a large increase. Once the UT1 errors were corrected, the charged-particle range estimate dropped back to an average of zero, as would be expected. To prevent this from happening again, the charged-particle range *a priori* sigma was reduced, after a quick study showed that it still performed well during earlier events. The interaction between the charged-particle model and UT1 is not well-understood, but is probably due to the effective station longitude offset (from the UT1 error) combining with tracking pass times that were not symmetric about the spacecraft's culmination time to produce a small Doppler bias, which the charged-particle model faithfully transformed into a range rate as long as the *a priori* sigmas allowed.

E. Spinning Spacecraft Effects

At this point the procession of measurement error discussion finally reaches the MER spacecraft itself. Like the Mars Pathfinder spacecraft before it, the MER spacecraft rotated at 2 rpm, and had a circularly polarized antenna with an offset from the spin axis. The removal of the sinusoidal signature from the antenna offset and the Doppler bias due to the polarized antenna is described in detail in Ref. 4, but apart from these deterministic effects there remains the possibility of an error in the magnitude of the Doppler bias. Before launch, a per-pass Doppler bias of 0.005 mm/sec (equivalent to 0.008 rpm) was included to account for this effect, although strictly speaking the Doppler bias should have been correlated on a many-day timescale since the spin rate error would have persisted. Post-launch experience with automatic spin-signature removal was convincing enough to remove the parameter from the filter, with the expectation that any remaining errors would behave as white noise and be covered in the Doppler weight.

F. Mars Ephemeris

Even though the Mars ephemeris is not related to an observation directly, the errors in the ephemeris behave like an observation error model, since they did not interact significantly with the trajectory until within the last 2 days. All the trajectory parameters of interest are Mars-relative, so the spacecraft state at the end of the data arc is mapped to the Mars entry interface (125 km altitude) and to Mars periapsis. Since there was not much sensitivity to the Mars ephemeris at even the TCM-6 data cutoff (~E-13 hours), the ephemeris errors were considered for all delivered solutions.

Of course, the relative insensitivity to the Mars ephemeris depends on ephemeris errors that are similar to the current-state errors at the end of the data arc. Since the Mars Odyssey spacecraft transitioned to mapping operations in February 2002, monthly Δ DOR observations were made of both Mars Odyssey and Mars Global Surveyor (which are tied to Mars to at least an order of magnitude better in position than the Earth-relative Δ DOR measurement accuracy). The resulting data were combined with range measurements to both spacecraft to produce an updated Mars ephemeris, which was internally delivered to the MER Navigation team as DE 410. The initial covariance (which had originally been provided with DE405+) used with this ephemeris had radial, downtrack, and normal sigmas of 17 meters, 271 meters, and 884 meters, respectively, in January 2004. However, an examination of the Δ DOR and range residuals used in the ephemeris fit strongly suggested that this covariance could be scaled down by at least a factor of two, which was the adopted value for baseline cases. The resulting 442-meter out-of-plane sigma is equivalent (at the Earth-Mars range at MER arrival) to the angular error from the quasar locations (described below) or to several Δ DOR measurements being statistically combined. Treating the ephemeris errors as consider parameters statistically combined (in an RSS sense) the mapped trajectory errors with the ephemeris for most of the approach, and only caused additional B-plane or EFPA uncertainty at the very end, when there was modest sensitivity to some of the parameters.^{§§§} Although included in the trajectory force model, the sensitivity to the Mars gravity field and mass errors was small, as is typically the case for hyperbolic approach trajectories, so errors for these parameters were not included in the filter.

G. Quasar Locations

The final parameters affecting the observations (in our discussion order) are the errors in the location of the quasars used as references in Δ DOR data. The pre-launch value of 2 nrad was retained throughout the mission, although the quasars in use were probably known to 1 nrad. As the quasars are not moving appreciably in an angular sense with respect to the Earth, the quasar location errors were always considered. Since the quasar location

^{§§§} Unfortunately, the very good knowledge of the MER spacecraft trajectory did not translate into another Mars ephemeris update, since the MER spacecraft did not end up tied to Mars at better than a few-hundred-meter level, instead of the few-meter level that would have resulted in a flyby. However, this wasn't the objective of this mission...

errors were assumed to be primarily due to data noise and quasar structure, there was no correlation modeled between any of the locations.^{****} Despite the small value of the quasar location error, this error and the Mars ephemeris error became the largest consider error sources, due to the large number of Δ DOR points and the low non-gravitational acceleration uncertainty level of the spacecraft.

H. Data Weights

As has been mentioned above, the main impetus behind the attention to detail in error modeling is the desire for consistency of solutions when one or more data types are withheld, or when other reasonable filter variations are exercised. Having concluded the discussion of observation error parameters, we now turn our attention to the proper modeling of the time-series of errors that represents the remaining noise in the data, and the relationship between the different data types. The filter treats the weight for each point as that point's element in a diagonal weighting matrix, thus assuming that all points are uncorrelated from each other in their weights.^{††††} However, since all of the estimated parameters have timescales much longer than the point spacing (one minute for Doppler, 12 minutes for range), and since all of the data types have diurnal periodicity, it is useful to develop a means for weighting each tracking pass (lasting up to 12 hours) with a weight derived from that pass's residual standard deviation. The Δ DOR data weighting also uses the pass as a consideration for weighting, but each pass is only two points separated by less than an hour. One final consideration for all data types is that a deweighting function based on elevation angle is applied to all data, such that the effective data weight under 20 degrees is much looser. Even though this elevation treatment was developed empirically in an earlier era when stochastic media modeling was not practiced, and consequently could stand to be reexamined, this method was used for the MER mission without any modification.

1. Doppler and range weights

Pre-launch data weights were 0.075 mm/sec for Doppler and 4 meters for range. Post-launch results quickly showed that the range was much better than 4 meters, which was a worst-case value for late in the MER-B trajectory. The Doppler weight was closer to the standard deviation, with the very best passes exhibiting values under 0.02 mm/sec. In mid-cruise (after ranging parameters were updated to account for the decreasing link performance), a typical range weight was 1 meter, and the original Doppler weight was retained. However, significant pass-to-pass variation was observed, due to the solar events noted above.

Doppler and range data are nearly continuous compared to any other model in the filter. Consequently, they can be sampled or auto-correlated at different timescales to determine the nature of the long-period noise. While this was not done for this mission, previous study has shown that tracking data noise follows a Kolmogorov power spectrum, with longer timescales having more noise than white phase or range noise would suggest. Since all forms of transmission media are turbulent phenomena, this is an expected result. The longest timescales (corresponding to the eddy generation scale) range from an hour for storm cells in the troposphere to many days for solar plasma. The correct statistical modeling for this is difficult to implement in an ideal filter, let alone one with operational limitations and a long development history, so an approximate approach is necessary. The adopted assumption is that Doppler and range data are primarily affected by the diurnal period, and anything longer is unimportant.^{****} If all the Doppler noise is due to a turbulent process, it can be shown that the relationship between the standard deviation of different timescales implies a weight that is proportional to the one-sixth power of the timescale ratio. Thus the 60-second Doppler data standard deviation would be multiplied by 3.36 to find the adopted weight, assuming one day as the longest relevant timescale. The range data (with a 12-minute sample time) is likewise multiplied by 2.22.

Both of these factors were applied to each pass of Doppler and range data, with additional constraints. As each new segment of tracking data came in, residuals were generated from the previous solution, with data editing being applied as necessary (which was rarely the case due to efficient pre-preprocessing). The mean, standard deviation, and maximum deviation (with respect to the mean) were all updated, and a file of weights generated. The maximum deviation divided by three replaced the standard deviation if it was larger, and a minimum weight of 0.05 mm/sec

^{****} The most accurate Mars approach trajectory estimates result when Δ DOR data is obtained from both Mars-orbiting spacecraft and Mars-approaching spacecraft with respect to the same quasar, which is easily possible for the last month or two. In this case the quasar error also cancels out for Mars-relative state knowledge, but the processing has to be done carefully to preserve this effect. For MER, the last Δ DOR observation of MGS or Odyssey was made several months before the MER-A arrival, so this technique could not be used, and wasn't necessary to meet the requirements.

^{††††} This is of course untrue for Doppler, since adjacent Doppler points share a common intervening phase measurement, but long-period noise makes this unimportant for timescales above tens of minutes.

^{****} Obviously for Mars orbiters, the relevant timescale would more likely be one orbit period, typically 2 hours.

and 14 centimeters was enforced as a lower limit. In the process a file with all of this information, as well as the pass start and stop times, was maintained. Due to the relatively low gain of the MER medium gain antenna, the range data (which used the same integration time for all passes during approach) clearly shows the significantly better performance of 70-meter antennas over the 34-meter antennas in the resulting residual standard deviation. However, the dominant range error in the model was still the per-pass range bias, so the change in range noise due to different tracking stations only served to shift the relative importance of the range and Doppler data within a pass in determining the angular position of the spacecraft.

2. Δ DOR weights

Pre-launch analysis produced a Δ DOR error budget with a value of 0.12 nsec (of relative delay), which was used for the covariance analysis, but this value was reconsidered during cruise. Despite the large, dense Δ DOR data set collected for the MER approaches, the total number of points (coming as they did at a maximum of two sessions per day, one on each of the two baselines) was not enough to develop independent statistics, or certainly not before both spacecraft had arrived at Mars. However, the Δ DOR had extensive pre-processing and analysis of similar measurements for Mars Odyssey and MGS. As mentioned above, each Δ DOR observation during approach consisted of two independent scan sequences, which were differenced into two points with a separation of about 45 minutes. The experience with the Mars orbiters had shown that the dominant Δ DOR error was instrumental biases with an expected magnitude of as much as 0.06 nsec, so this was taken as a 1 sigma value and used as the minimum for each session. With two points in a session being treated by the filter as uncorrelated, this required 0.085 nsec as the weight for each point. Since there wasn't a session with only 1 point in the last 2 months, 0.085 was adopted as the Δ DOR weight.^{§§§§}

I. Spacecraft Modeling

Whereas the gravitational and relativistic acceleration on a spacecraft in interplanetary space are known almost perfectly (such that modeling the remaining errors is unimportant), the non-gravitational forces on the spacecraft are often poorly known, and are almost always the limiting factor for trajectory prediction. The MER spacecraft had very well-behaved non-gravitational accelerations, due to being a spinning spacecraft and having balanced thrusters, but the modeling of these parameters still made a big difference in the final result. Specifically, the error modeling of the solar pressure and of the propulsive events (attitude changes and trajectory correction maneuvers (TCMs)) is described below. With the exception of the TCMs, all non-gravitational error models are included all the way to the target by propagating constant errors or continuing to update stochastic parameters, instead of stopping error models at the end of the data arc.

1. Solar pressure radiation

Pre-launch modeling of the solar pressure used flat plates and cylinders to model the cruise stage and approximate the lit region of the backshell. The area of each component was estimated with a 5 percent *a priori* sigma, and each reflection coefficient was estimated with an *a priori* sigma of 10 percent of its value. An additional spherical acceleration was estimated stochastically with a 1-day update time, a 10-day correlation time, and a process noise of 2×10^{-12} km/sec² to account for solar pressure errors in all three directions. However, it was recognized that both the nominal model and its errors should be updated when time permitted, which unsurprisingly turned out to be after launch.

The details of the updated solar pressure model are




Figure 4: Solar aspect angle history for MER-A and MER-B.

^{§§§§} A more complicated weighting scheme could be imagined, using the residual difference of the two points of each session as input, but experiments along these lines were inconclusive. On two occasions, Δ DOR observations were made simultaneously with four stations (two at each complex). This could help to separate station biases from troposphere/ionosphere errors, but would have had to have been done more often to develop reliable statistics.

given in Ref. 4. Since the average surface normal vector for lit areas always fell in the Sun-spin axis plane, there was no solar pressure force out of that plane, making the acceleration modeling a 2-dimensional problem.^{*****} The biggest uncertainty in the model was the overall radiation balance, since any heat that did not depart through the radiators (with no net effect due to symmetry) was radiated off of the top of the cruise stage instead of off the heat shield (which was always in shadow, and hence at or below -100 degrees C), producing an additional force beyond what would be expected from the material properties of the solar arrays. The solar pressure model accounted for this effect by increasing the diffuse reflection for the flat plate representing the top of the cruise stage to a physically unreasonable value, but one which did a good job of modeling the force for Sun-spin axis (or solar aspect) angles of about 35 degrees. Since the solar aspect angle determined the total lit area of the spacecraft, and hence the total heat load, the cruise stage re-radiation would be expected to depend on the solar aspect angle. During Mars approach, the solar aspect angle varied between 25 and 45 degrees, changing about half a degree per day in between turns, which ranged from 5 to 18 degrees and occurred every 8 to 50 days, as shown in Fig. 4.

Consequently, the solar pressure error model included a stochastic diffuse reflection for the flat plate normal to the spin axis (representing the top of the cruise stage), with a process noise of 0.1 (as a coefficient), an update time of 1 day, and a 7-day correlation time. This parameter's correlation was broken at each turn, assuming that the errors would not necessarily be similar at different solar aspect angles. In addition, the specular and diffuse coefficients for each of the three components of the model (a flat plate, a cylinder for the sides of the cruise stage and aeroshell edge, and a tabular function for the lit area of the backshell) were estimated as constants with *a priori* sigmas of 0.1 (or 10 percent of the maximum physical coefficient value of one).^{††††} In practice, the stochastic diffuse coefficient did not vary appreciably across turns, except for early in the MER-B approach where the solar aspect angle reached almost 45 degrees, and even there the difference was less than half of the process noise. From this, it was concluded that the solar pressure model was very accurate at the approach geometry, which contributed significantly to the consistency of the approach solutions.

2. Attitude changes

Although spinning spacecraft can go a long time without using their propulsion system, changing the spacecraft spin axis requires using thrusters and is generally required to maintain either power or communications within reasonable bounds. In the case of the MER spacecraft, the attitude changes (also referred to as ACS events in Tables 1 and 2) were required to maintain communications on the medium gain antenna, but the relatively low gain of the antenna allowed the attitude changes to be correspondingly infrequent, even for a spinning spacecraft. The thruster configuration (described in more detail in Ref. 1) allows turns to be made in a balanced mode, but concern over the degree of imbalance due to misalignment, plume impingement, and impulse mismatch caused fairly large spherical uncertainties of 3 mm/sec to be assumed before launch. This was perhaps understandable due to recent unpleasant experience with spacecraft with unbalanced thrusters, but it still constituted the most significant factor in under-predicting the actual performance (in the sense of causing the predicted delivery errors to be too large). Regardless, post-launch experience (including the ACS-NAV calibration activity described in Ref. 4) quickly reduced these uncertainties to 0.5 mm/sec for MER-A and 1 mm/sec for MER-B.^{††††}

Even at this level, the attitude turns were a dominant component of the delivery errors. Further analysis of the ACS-NAV calibration⁴ showed that the unbalanced components for typical turns were 0.02 mm/sec or less except for the spin-axis component of MER-B turns, which was 0.1 mm/sec. This led to the adoption of 0.05 mm/sec as the *a priori* sigma for all components except MER-B axial, which had an *a priori* sigma of 0.1 mm/sec and a nominal value of 0.1 mm/sec (for typical 10 degree turns). At the same time, the number of attitude turns had been cut roughly in half due to a trade between the amount of commanding activity on the spacecraft and the telecommunications performance. Although the ACS-NAV calibration had been designed for the original 5-degree turns, the results for 10 degrees were not found to be twice as large, so the net result of the decrease in the number of turns was increased trajectory knowledge. For the delivery at entry, the most important turn occurred right after TCM-4 on each spacecraft, which was placed 8 days before entry. The effect of the reduction in the ΔV uncertainty, and to a lesser extent the reduction in the number of turns, produced B-plane delivery accuracies that were

^{*****} Of course, re-radiation from surfaces passing in and out of shadow could have caused an out-of-plane force, but the expected temperature variations made this too small to notice, and no evidence for such a force was seen in the ΔDOR data.

^{††††} In the JPL Orbit Determination software, the specular and diffuse coefficient values are divided by two and three, respectively, before being input to the program, which is why Tables 1 and 2 show 0.05 and 0.033 for the sigmas of these parameters.

^{††††} MER-B was actually the first vehicle built, and the thruster alignment was known to be slightly worse than for MER-A before launch. This proved true in flight, though not to a significant degree.

dominated by the nominal TCM execution errors, and surface delivery accuracies that were dominated by atmospheric uncertainty.

3. TCM-4

The accuracy of the solar pressure model and the earlier maneuvers reduced the size of TCM-A4 (TCM-4 on MER-A) to a single pulse of 3.4 seconds duration, with a ΔV of 2.5 cm/sec. The pre-launch maneuver execution error model included spherical fixed errors of 2 mm/sec (1 sigma), statistically combined with spherical errors of 1.67 percent of the magnitude of the total inertial ΔV ⁵. These numbers were about as conservative as the original attitude turn assumptions, especially for small maneuvers, and the fixed errors were particularly burdensome. Some improvement might have been assumed by differentiating between control errors and *post facto* knowledge errors (when telemetry from the maneuver is available), but this was not done for MER.

Consequently, TCM-A4 was modeled with a 1.67% magnitude error, and a 0.5 degree pointing error, based on discussions with ACS engineers. A 3-second start time uncertainty was also included, which proved to be much larger than necessary, but harmless to the delivery results. As will be discussed below, a variety of values around these baseline values were also investigated, and the data eventually determined the post-maneuver trajectory somewhat independently from the maneuver assumptions. However, solutions with the tighter *a priori* sigmas arrived at the right answer for TCM-A4 more quickly.

For both spacecraft, the estimation of TCM-4 was slightly corrupted by the attitude turn immediately following the maneuver, but the small values of the turn uncertainty reduced this effect. If TCM-5 had been required on either spacecraft, it would have benefited from not having the post-maneuver turn, but suffered from having only 3 Δ DOR points afterwards before entry.

For MER-B, operational priorities had resulted in the cancellation of TCM-B3, leaving the spacecraft with a larger TCM-B4 (though still smaller than any maneuver other than TCM-A4), with an inertial magnitude of just over 20 cm/sec. The resulting maneuver had an axial component of 8.3 cm/sec (20 seconds of firing 2 thrusters), and a lateral component of 20 cm/sec (just under 2 pulses, with on times of 5 seconds and 4.991 seconds). The original maneuver execution error model did not specify any different treatment for a vector mode maneuver, but clearly consideration must be given to the existence of two different sorts of events. Consequently, the same nominal *a priori* sigmas were applied (1.67 percent magnitude, 0.5 degrees pointing in each direction), except that the segment start time sigmas were reduced to 1 second. With two ΔV events occurring this close together, any estimate of the non-line-of-sight components is necessarily highly correlated. As before, a variety of values around these nominal ones were used, with the final results showing corrections of more than one sigma (but less than two), so looser *a priori* sigmas would have been justified. Nonetheless, the post-TCM-B4 solution converged before the TCM-B5 solution was delivered.

III. Filter Variations

The definition of a baseline case is no more than a starting point in the quest to discover the truth about the spacecraft's orbit. Over the course of the approaches of the two spacecraft, the baseline case (as defined by the solution delivered to the rest of the Navigation team) changed 4 times, in response to discoveries and new revelations about the orbit determination process. However, the maintenance of a baseline case and variations around it proved a useful discipline in managing the solution process, including describing it to others. The discussion of filter variations will proceed from solution case descriptions to case organization and implementation, and conclude with event response.

A. Case Descriptions

In the parlance of the MER Navigation team, a "case" is a particular set of tracking data and filter implementation, ideally independent of the tracking data cutoff. Consequently, each case has a succession of instances, each with a later data cut off time, all of which can be evaluated for appropriate solution behavior. The majority of the cases used were not reasonable as delivery candidates, but were intended to illustrate one facet of the behavior of the family of baseline cases, and may be thought of as excursions along one dimension of a hyperspace, which turns out to have (insofar as the analogy holds) four dimensions in this implementation.

The first two dimensions have to do with the tracking data included in the solution. While all of the cases advance their endpoints simultaneously as they are updated with new data, their starting points are staggered to develop a range of arc lengths. Table 3 shows the start times for the 7 MER-A and MER-B data arcs, along with the

Table 3: Start Times for Orbit Solution Arcs

a) MER-A

b) MER-B

Arc Name	Arc Length (days) prior to		Reasoning
	TCM-A4	Entry	
Very Long	66	74	SAA above 30 deg. 2 ΔDOR per baseline before TCM-A3
Long	43	51	After TCM-A3
Medium	32	40	After ACS-A8
Short	16	24	After ACS-A10
Very Short	7	15	7 day moving window starting arc before TCM-A4
Post TCM-A4	-	8	After TCM-A4
Entry	-	2	Allow entry Doppler to drive solution

a)

Arc Name	Arc Length (days) prior to		Reasoning
	TCM-B4	Entry	
Very Long	88	96	SAA above 30 deg. 2 ΔDOR per baseline before TCM-B3
Long	45	53	After TCM-B3
Medium	27	35	After ACS-B8
Short	13	21	After ACS-B10
Very Short	7	15	7 day moving window starting arc before TCM-B4
Post TCM-B4	-	8	After TCM-B4
Entry	-	2	Allow entry Doppler to drive solution

b)

rationale for that particular choice. Rather than maintaining a moving window for short-arc solutions, the shorter arcs were added as soon as at least 1-2 days of tracking data was available after their starting points. The utility of a range of arc lengths is their varied response to different error scenarios. For instance, the response to an acceleration modeling error (combined with the rather low process noise on the spacecraft state) would have been poor fits on the longer data arcs, with the shorter arcs jumping more quickly to the new solution. As it turned out, the acceleration modeling was essentially correct throughout, and so the longer arcs consistently displayed both good fits and slightly better delivery statistics, due to their larger amount of data. Consequently, although any solution with 3-4 weeks of data could have been delivered, the long or very long arcs were uniformly used for all deliveries.^{§§§§§} Figure 5 shows the baseline case for 7 data arcs for MER-A for a data cutoff on January 3, 2004, 03:45 UTC. Except for a 14-percent growth in the semi-major axis of the ellipse, there is little to choose between the 4 longest arcs, and all the arcs are consistent to well within one sigma.




Figure 5. B-plane plot of solutions from 7 data arcs for MER-A on January 3, 2004

^{§§§§§} The only difficulty with the very long arc was the processing time, which became onerous for both arcs, but less so for the long arc.

Mars B-plane (Mars Equatorial of Date)

The plot shows the Mars B-plane with the horizontal axis labeled $B \cdot T \text{ (km)}$ ranging from 6680 to 6720, and the vertical axis labeled $B \cdot R \text{ (km)}$ ranging from 2540 to 2580. Several regions and orbits are depicted:

- Range, ΔDOR** : A large cyan ellipse representing the range measurement region.
- Doppler, ΔDOR** : A magenta ellipse representing the Doppler measurement region.
- Dep, Range, NS- ΔDOR** : A green ellipse representing the combined Doppler and range measurement region for the NS- ΔDOR configuration.
- Doppler, Range**: A black ellipse representing the combined Doppler and range measurement region.
- Doppler-only**: A blue ellipse representing the Doppler-only measurement region.
- Dep, Range, EW- ΔDOR** : A blue ellipse representing the combined Doppler and range measurement region for the EW- ΔDOR configuration.

Small crosses mark the centers of the ellipses, and a line connects the centers of the 'Dep, Range, NS- ΔDOR ' and 'Dep, Range, EW- ΔDOR ' regions.

Figure 6. B-plane plot of 6 data type combinations for MER-A on December 25, 2003

Table 4. Orbit solution case list matrix as of December 26, 2003.

Class:		+00	+01	+02	+03	+04	+05	+06	+12	+13	+14
Data Series *	Filter Setup:	none	Baseline	Consider Media/EOP	Consider Media/EOP	Baseline	RNG bias per DSS	Baseline	RNG bias per DSS, Consider Media/EOP	RNG bias per DSS, Consider Media/EOP	RNG bias per DSS
	Data Weighting:	none	Weight By Pass (WBP)	Global Wgt, Deweight No Cals	WBP	Global Weight	WBP	WBP, DDOR conserv.	Global Wgt, Deweight No Cals	WBP	Global Weight
	100 F2, SRA, DDOR	0	1	0	1	1	1	1	0	0	
200 F2, SRA	0	1	0	0	0	0	1	0	1	1	
300 F2, DDOR	0	1	0	0	0	0	0	0	0	0	
400 F2	0	1	0	0	0	0	0	0	0	0	
500 F2, SRA, EW-only	0	1	0	0	0	0	0	0	0	0	
600 F2, SRA, NS-only	0	1	0	0	0	0	0	0	0	0	
700 SRA, DDOR	0	1	0	0	0	0	0	0	0	0	
Tally of Cases		0	7	7	8	9	11	12	13	14	1

Class:		+20	+21	+22	+23	+24	+25	+26	+27	+28	+29
Data Series *	Filter Setup:	Accel nom	Loose SRP	Loose ACS	Loose DOPRNG	Baseline	Baseline	Baseline	Baseline	Tight MNVR	Est Mars Eph, 1+DE405-sigma
	Data Weighting:	WBP	WBP	WBP	WRP, Loose DOPRNG	WRP, PT since 12/17**	WRP, DDOR PT since 12/17**	WBP Tight	WBP, High Elev	WBP	WBP
	100 F2, SRA, DDOR	1	1	1	0	1	1	1	1	0	0
200 F2, SRA	0	0	0	0	1	1	0	0	1	0	0
Tally of Cases		16	17	18	19	21	22	23	24	24	2

Class:		+30	+31			+35			
Data Series *	Filter Setup:	Charge particle delay model	Random walk SRA: wht noise. F2 bias						
	Data Weighting:	WBP	WBP			Baseline WBP, F2/SRA PT since 12/17**			
	100 F2, SRA, DDOR	1	1			1			
200 F2, SRA	1	1			0				
Tally of Cases		26	28			29			

* Add series to class to get case number.	B-plane history mapping No history mapping
** Data are not deleted, just de-weighted.	

Figure 6 shows the B-plane covariance for the six combinations without all the data for the MER-A very long arc with a December 25, 20:47 UTC data cutoff. The Doppler-only case is consistent with the others, but much larger, as expected for a data set that lacks direct position measurements. The range and Δ DOR case is indistinguishable from the case with all the data, which is consistent with the relatively low information content in the Doppler data. In between these extremes, the Doppler and range case is consistent with the others, and the Doppler and Δ DOR case shows that the B-plane is not aligned with the geocentric plane-of-the-sky. The Doppler, range, and EW- Δ DOR ellipse is larger than the Doppler and range case in the $B \bullet R$ direction due to consider effects, and the powerful effect of the intersection of the EW- Δ DOR and the NS- Δ DOR is evident.

The third and fourth dimensions are a 2-dimensional matrix filter variations, where individual cases have been grouped by either likelihood, or by a shift in the baseline case. In Table 4, cases 01 through 06 are all small variations around the baseline case 01. Cases 12-14 include range biases per station, but otherwise duplicate cases 02-04. Cases 20-29 explore less-reasonable dynamics and data weighting variations, or pass through the data at the end, and cases 30 and 31 address the solar plasma problems with two different models. By the time of MER-B final approach, all of the cases in the 30s and 40s had been filled in, as shown in Table 5.

Table 5. Orbit solution case list matrix as of January 20, 2004.

Data Series*	Class:										# OD Cases***
	+50	+01	+02	+03	+04	+05	+06	+12	+13	+18	
Filter Setup:	Charge particle delay, No planet Eph. Considered	Baseline	Consider Media/EOP	Consider Media/EOP	Baseline	RNG bias per DSS	Baseline	RNG bias per DSS, consider Media/EOP	RNG bias per DSS, consider Media/EOP	Charge particle delay, Tight TCM-B4, 14-day const. DIFF01	
Data Weighting:	WBP	Weight By Pass (WBP)	Global Wgt. Deweight No Cals	WBP	Global Weight	WBP	WBP, DDOR conserv.	Global Wgt. Deweight No Cals	WBP	WBP	
100 F2, SRA, DDOR	1	1	0	0	0	0	0	0	0	0	7
300 F2, DDOR	0	0	0	0	0	0	0	0	0	0	32
400 F2	0	0	0	0	0	0	0	0	0	0	33
700 SRA, DDOR	0	1	0	0	0	0	0	0	0	0	36
Tally of Cases	1	5	5	5	5	5	5	5	5	6	
Data Series*	Class:										# OD Cases***
	+20	+21	+22	+23	+24	+25	+26	+27	+28	+29	
Filter Setup:	Accel nom	Loose SRP	Loose ACS	Loose DOPRNG	Baseline	Baseline	Baseline	Baseline	Charge particle delay, Tight TCM-B4	Est Mars Eph, 1*DE405+ sigma	
Data Weighting:	WBP	WBP	WBP	WBP, PT since 12/27**	WBP, PT since 12/27**	WBP Tight	WBP, High Elev	WBP	WBP	WBP	
100 F2, SRA, DDOR	0	0	0	0	0	0	0	0	1	1	6
Tally of Cases	6	6	6	7	7	7	7	7	8	9	
Data Series*	Class:										# OD Cases***
	+30	+31	+32	+33	+34	+35	+36	+37	+38	+39	
Filter Setup:	Charge particle delay	Random walk SRA: wht noise, F2 bias	Charge particle delay, RNG bias per DSS	Charge particle delay, consider Media/EOP	Charge particle delay, Baseline	Charge particle delay, Baseline	Charge particle delay, Baseline	Charge particle delay, RNG bias per DSS, 14-day const. DIFF01	Charge particle delay, Loose TCM-B4	Charge particle delay, Est Eph.	
Data Weighting:	WBP	WBP	WBP	WBP	Global Weight	WBP, PT since Jan 17**	WBP, DDOR conserv.	WBP	WBP	WBP	
100 F2, SRA, DDOR	1	1	1	1	1	1	1	1	1	1	16
200 F2, SRA	0	0	0	0	0	0	0	0	0	0	28
500 F2, SRA, EW-only	0	0	0	0	0	0	0	0	0	0	34
600 F2, SRA, NS-only	0	0	0	0	0	0	0	0	0	0	35
Tally of Cases	13	15	17	18	19	20	21	22	23	24	
Data Series*	Class:										# OD Cases***
	+40	+41	+42	+43	+44	+45	+46	+47	+48	+49	
Filter Setup:	Charge particle delay, Accel nom	Charge particle delay, Loose SRP	Charge particle delay, Loose ACS	Charge particle delay, No DIFF01 stochastic	Charge particle delay, Baseline	Charge particle delay, Baseline	Charge particle delay, Baseline	Charge particle delay, Baseline	Charge particle delay, Very Loose TCM-B4	Charge particle delay, Very Loose TCM-B4, 14-day const. DIFF01	
Data Weighting:	WBP	WBP	WBP	WBP	WBP, PT since Jan 17**	WBP, DDOR PT since Jan 17**	WBP Tight	WBP, High Elev	WBP	WBP	
100 F2, SRA, DDOR	1	1	1	1	1	1	1	1	1	1	26
200 F2, SRA	0	0	0	0	0	0	0	0	0	0	37
Tally of Cases	25	26	27	29	31	32	33	34	35	36	

* Add series to class to get case number
 ** Data are not deleted, just de-weighted.
 ***# OD Cases* help analysts keep track of where filter loop is in its run.

Abbreviated Set of Cases
 Completes Full Set of Cases
 Not Run on Post TCM-B4 arcs

Generally, each new filter case was formed by varying a single element of either the baseline case or a potential baseline case. An early exception to this was cases 01-04, which filled out the matrix of estimate vs. consider media and Earth orientation parameters (EOP), and the weight-by-pass (WBP) vs. global weighting. Since the baseline case (estimating media and EOP, using WBP) proved most effective, the other cases were not generally retained. The first three cases in the 20s examined adding stochastic accelerations at a fairly unreasonable level (3×10^{-12} km/sec²), loose solar pressure (45-50 percent coefficient sigmas), or loose attitude (ACS) turns (0.5 mm/sec spherical uncertainty). To varying degrees, these all have the effect of shortening the data arc, and are primarily useful for model identification in the event of an acceleration modeling error. Cases 24, 25, and 35 examine the prediction capability for earlier data, now that later data is available, with a pass-through date that was periodically moved to keep the pass-through duration at about a week. Cases 23, 26, and 27 vary the relative data weights between the Δ DOR and the Doppler and range data, as a less extreme version of entirely deleting one data type.

For the same arc and data cut-off as Fig. 6, Fig. 7 shows the baseline case against cases with either looser dynamics (120, 122), earlier data cutoffs (124, 125), or different data weighting (126). The inconsistency of case 125 is probably due to the growing UT1 error mentioned above. The tighter Doppler and range




Figure 7. B-plane plot of cases with varying dynamic error models.

weighting in case 126 is trying to move the solution in the same direction, but is constrained by the last week of Δ DOR data. Not surprisingly, the two cases with loose dynamics move in the same direction, but not to an inconsistent degree. *****

B. Case Organization and Implementation

The organization of the filter cases has already been displayed in Tables 4 and 5 in the preceding discussion, but not explicitly stated. Ideally, a recursive hash (whose elements may be references to another hash) of arbitrary depth would be used to allow any combination of cases to be described, but this did not fit well with the existing software (as described below) or, upon reflection, with human discussion of the cases. The approach that was adopted has essentially 3 indices: a leading letter to indicate arc length, and, following a period, a 3-digit number, where the first digit (in the hundreds place) indicates the data type set. While cases were grouped at the tens level, no explicit requirement was made that this be done. So case “A.101” is a very-long arc case with all data and the original baseline filter, and case “C.232” is a medium arc case with Doppler and range data which uses the charged-particle model with range biases estimated per station. In discussion, the 3-digit numbers were used most frequently, with the letters being replaced by words describing the arc or dropped altogether.

Tables 4 and 5 also show which cases were run under what circumstances. Although the cases were run with a high degree of automation (as described below), the longer arcs could easily take an hour or more to run 40+ cases, even after dividing the arcs across different computers. Consequently, a hierarchy of cases was developed with a color code, such that the most important ones could be run first, and then the others filled in, as shown in Table 5.

Regardless of the beauty of the case list, it would have taken a small army of analysts to make the required runs if any manual input was needed for each case. This approach was actually used as recently as 2001 for the Mars Odyssey approach^{†††††}, but clearly needed improvement. A program was developed to apply a file of variations to a baseline case automatically, and store the results both in a database (for the more important summary data) and in directory structures created automatically. Although the maintenance of both the case list file at a working level, and the many files it pointed to, was not simple, the benefits were enormous, and small changes were not hard to make once the initial setup had been developed. As is often the case, a small suite of scripts was also developed around the main programs to prepare case list input files and further the automation. For example, one such script controlled which case colors (such as shown on Table 5) were run, without having to resort to hand editing the case list file, and another updated the epochs reflecting the last media calibrations and tracking data received.

Once the filter cases were generated for the current data cutoff, solutions needed to be compared against each other. Although a variety of comparisons were made occasionally, the B-plane ellipses were by far the most common and useful comparison. Many of the restrictions of the case naming nomenclature were imposed by the need to be able to efficiently select particular cases from the database for display. Figure 8 shows the graphical user

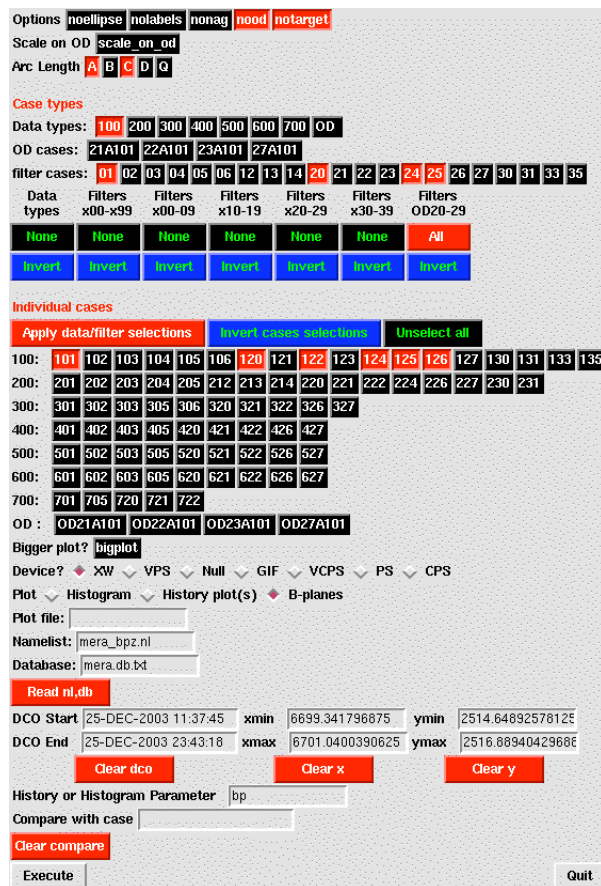


Figure 8. Sample “xbpz” graphical user interface for B-plane plot control.

***** The shorter data arcs also tend to be offset up in B•R for this data cutoff.

††††† Not that the number of analysts was larger, but many fewer solutions were generated less often.

interface (known as ‘xbpz’) developed^{*****} to make this possible. The intent of the selection options is to be able to control the solution set to be displayed with a minimum of key or mouse strokes. The four dimensions of variation are reflected in the arc length, data type, and filter case areas. Once a solution set and any scaling or display options had been selected on the xbpz GUI, the “Execute” button generated a B-plane plot, as shown in Fig. 9. Until the B-plane plot was dismissed, only limited rescaling of the plot was possible. In Fig. 9 the blue circle is the MER-A TCM-4 delivery, and the green arc (which looks like a line) is the target B-magnitude (corresponding with the EFPA of -11.5 degrees). The magenta ellipses are previous delivered solutions, with the largest one having been made before the more realistic *a priori* sigmas were adopted for TCM parameters.

C. Event Response

The evolution of the case list is already evident from Table 5, which shows the final case list for MER-B. The development of the charged-particle model made the original cases in the zeros, 10s, and 20s obsolete, so few of them were still run. The new baseline used the charged-particle model (case 30), so some cases in the zeros were duplicated as variations around case 30 (e.g., Case 03 to 33, 04 to 34, 05 to 32 since 35 was already defined as a variation to 25). Likewise, cases in the 20s were duplicated in the 40s, starting from case 30 (e.g., Case 20-22 to 40-42, etc.). Regardless of the irrelevance of a previous case, no case numbers were recycled, to allow effective comparison of cases for different data cutoffs.

The biggest event in the late approach of both spacecraft was TCM-4. The databases were kept separate across the TCM-4 boundary, partly to reduce their size, but also because the B-plane shift (particularly for MER-B) made comparison across the TCM less useful. The deterministic ΔV of the TCM was modeled in all cases (for arcs starting beforehand) and the estimated parameters were updated. TCM-4 also marked the introduction of variations in the statistical modeling of the maneuver. Cases 28, 30 (the baseline case), 38, and 48 covered a spectrum of constraints from tight to very loose, as shown in Table 6 (for TCM-B4, but TCM-A4 sigmas were the same). As more data was obtained after the maneuver, all of these cases converged.

For TCM-B4 on MER-B, the post-maneuver stochastic solar pressure estimate was out of family with other estimates at similar solar aspect angles. It turned out that modeling errors above the one sigma level in TCM-B4 were aliasing into that stochastic parameter, but clearly cases were needed to prevent this from happening. In response, case 37 was added as a variation to case 32 which stopped updating the stochastic solar pressure for the last 7 days before TCM-B4 and thereafter. In addition, case 43 was added as a variation to case 30 which did not estimate the stochastic solar pressure at all (not unreasonable due to the statistically small corrections that had been observed). Both of these cases helped to force the right movement in the TCM-B4 parameters, and by the time of the TCM-B5 data cutoff, the original solar pressure treatment had converged with the new cases. Nevertheless, several solutions were delivered from case 137 following TCM-B4.




Figure 9. Sample B-plane plot generated by xbpz

Table 6. TCM-B4 error assumptions for filter case variations.

	A Priori value	A Priori Sigma			
		28 Tight	30 Baseline	38 Loose	48 Very Loose
Axial Thrust (N)	4.3722	0.073	0.073	0.109	0.219
Axial RA (deg.)	-160.8	0.10	0.50	1.00	2.00
Axial DEC (deg.)	-6.93	0.10	0.50	1.00	2.00
Lateral Thrust (N)	2.7965	0.047	0.047	0.070	0.140
Lateral RA (deg.)	-132.15	0.10	0.50	1.00	2.00
Lateral DEC (deg.)	-8.35	0.10	0.50	1.00	2.00

^{*****} To be fair to Mars Odyssey navigation, development of this tool was started during Odyssey’s Mars approach, although significant additional development was required to arrive at the current form.

IV. Statistical Tests and Presentation

The vast amount of data generated by running the solutions described above several times a day immediately presents the challenge of finding the nuggets of important information amongst the mounds of dross (results too similar to be remarkable). Since this was the first time that such an automated approach had been implemented, the demonstrated effectiveness of finding significant features in the data still leaves some room for improvement, but several useful tools were developed before the MER approaches to assist in this task. The process of evaluating the current status of the orbit determination began with statistical tests (of consistency between solutions and data content compared to *a priori* knowledge), followed by the automatic generation of history and data statistics plots, and concluding with an organized discussion of the results, aided by an automatically-generated report in HTML format.

A. Statistical Tests

A common test to quickly find the significant differences between two solutions is to divide the parameter estimate differences by the larger of the two parameter sigmas, and order the results from largest to smallest. JPL navigation has used a tool to perform this test for almost a decade, and found it a useful way to see what parameters are significantly different. However, this is only a succession of one-dimensional slices through a multidimensional covariance, and while parameter differences are an indication of solution differences, the actual quantities of interest (such as the B-plane or EFPA) may not differ very much.

For MER, a more rigorous test was developed to measure the consistency of solutions in multiple dimensions (although only the two B-plane dimensions were used). This method finds the smallest common scale factor that produces a point where two scaled ellipses are tangent. The common scaling enforces common likelihood for the tangent point with respect to each ellipse. Smaller scale factors are an indication of more consistent solutions. A detailed description of the mathematics behind this method is given in Appendix A.

Within each data arc, all of the solutions for each data cutoff had their mutual solution consistency factors calculated. The resulting matrix went through some iteration before arriving at a color-coded plot format to present the results. Whereas MER-A made do with just two colors for solution consistency, with a threshold of 0.8 for transitioning from red to blue, by MER-B approach the tool had been modified to use 5 colors, with thresholds at 0.2, 0.5, 0.8, and 1.2, as shown in Fig. 10. The value of the plot was to indicate which solution families were not agreeing, and also to help indicate which arcs were in the best internal agreement for delivery purposes. A more complete approach would have been to generate the solution consistency factor across all arcs for a given data cutoff, but this was not practical with the presentation tools available, not to mention the computer resources needed to perform the calculation (which obviously grows in size as N^2).

Another test considered the ratios of a parameter's correction and pre- and post-fit sigmas. The correction divided by the post-fit sigma determines the significance of the correction, while both the correction and the post-fit sigma divided by the pre-fit (*a priori*) sigma measures the importance and accuracy of the pre-fit sigma and nominal parameter value. These three ratios were computed for all of the estimated parameters (whether constant or stochastic), but were each evaluated only against a single threshold, and presented as text. While this method can currently be used as an input for backup analysis, it would need more effective ordering and reporting of large values to be more useful.




Figure 10. Sample solution consistency plot

B. Display Generation

The display tools used for MER orbit determination tried to strike a balance between making displays available on request, and pre-generating plots for the web page. Generally, plots of parameters within a single solution were only generated on request. An important example is the display of stochastic estimates and uncertainties, as shown in Fig. 11. Since this figure shows the range bias history, the tracking stations have been identified using labels and colors, but all other stochastic parameters used only one color. Pre-fit and post-fit tracking data residuals are




Figure 11. Range bias estimate history for MER-B.

another example of a plot generated on request. Examples of plots of Doppler residuals can be seen in Ref. 2 on page 9. In addition, the predicted post-fit residuals of a withheld data type can reveal the modeling errors in solutions based on the rest of the data. An example of Δ DOR pass-through residuals (in units of nsec) for a Doppler and range solution is shown in Fig. 12. Whereas the East-West points (indicated with a '6') are flat, the slope in the North-South points (indicated with a '4') show that there remains an out-of-plane velocity error, with the position error fortuitously vanishing at the end. Range pass-throughs of Doppler and Δ DOR solutions can also be useful to examine.




Figure 12. Δ DOR pass-through residuals for a Doppler and range solution.

One useful method for solution evaluation is displaying the history of that particular solution strategy. This can be done either within the same filter run, by saving results periodically, or by saving earlier final results from the same solution strategy for later display. Figure 13 is an example of the latter for a case with loose TCM-B4 constraints, and shows how the estimates converged in the days after the TCM (which was at Entry-8 days). These plots were generated for all current cases automatically for use with the web page.




Figure 13. Sample EFPA history plot.

After looking at the individual cases, plots were generated across all the cases in a directory (comprising one arc length/data cutoff). Figure 14 shows an example of plots comparing one data type's bias and sigma across all cases in a directory. These plots were also generated for cases that did not fit the particular data type (although that is not shown here), which was useful to evaluate the residual consistency for withheld data types.




Figure 14. Sample Δ DOR residual bias and sigma for cases fitting Δ DOR data.

Although the B-plane plotting tool shown in section III is another example of a tool used on demand, it was also used to pre-generate a defined list of B-plane plots, both of the solutions in a directory and for subsets of cases across the same data cutoff (but different arc lengths), for use with the web page. Just as for the other pre-generated plots, keeping the files specifying which plots to generate up-to-date (with the identity of the current baseline case, etc.) was sometimes a challenge.

C. Organized Discussion

The results of the statistical tests and the pre-generated plots were all brought together by a script (known as 'NavReport') that built and maintained a web page with links to most items of interest. A typical NavReport page is shown in Fig. 15. The frame on the left gave access to residual plots (generated as needed by starting another program), pre-generated B-plane plots, solution histories, solution consistency plots, solution summary information, and older versions of the preceding list for previous data cutoffs. The frame on the right was re-generated for each of the categories of information, and scrolled separately. More detailed information on each of the arcs was available through the arc name links at the top of the page. Although only visible in the flight operations environment, this web page was an easy way of making a summary of the current orbit determination status available to the rest of the Navigation team, as well as to orbit determination and navigation experts in the Navigation and Mission Design section.

Over the past several years, the realization of the difficulties faced by all Navigation teams in having time to consider all the relevant aspects of navigation operations led to the development of the Navigation Advisory Group (NAG), an informal group of experts^{#####} in the Navigation and Mission Design section. In addition to the web page mentioned above, an important means for interacting with the NAG was a meeting (called the 'Daily Show' as well as 'the NAG'), held nearly daily (usually at 4 pm) for the last several weeks of each spacecraft's approach. The challenge for this sort of interaction is to keep the extra-team experts sufficiently informed about the orbit determination and navigation status to allow them to contribute useful suggestions and insight, while still maintaining the ability to move quickly during time-constrained operational scenarios (such as final maneuver design). The Daily Show meetings accomplished this by providing a forum for interaction at a time in between

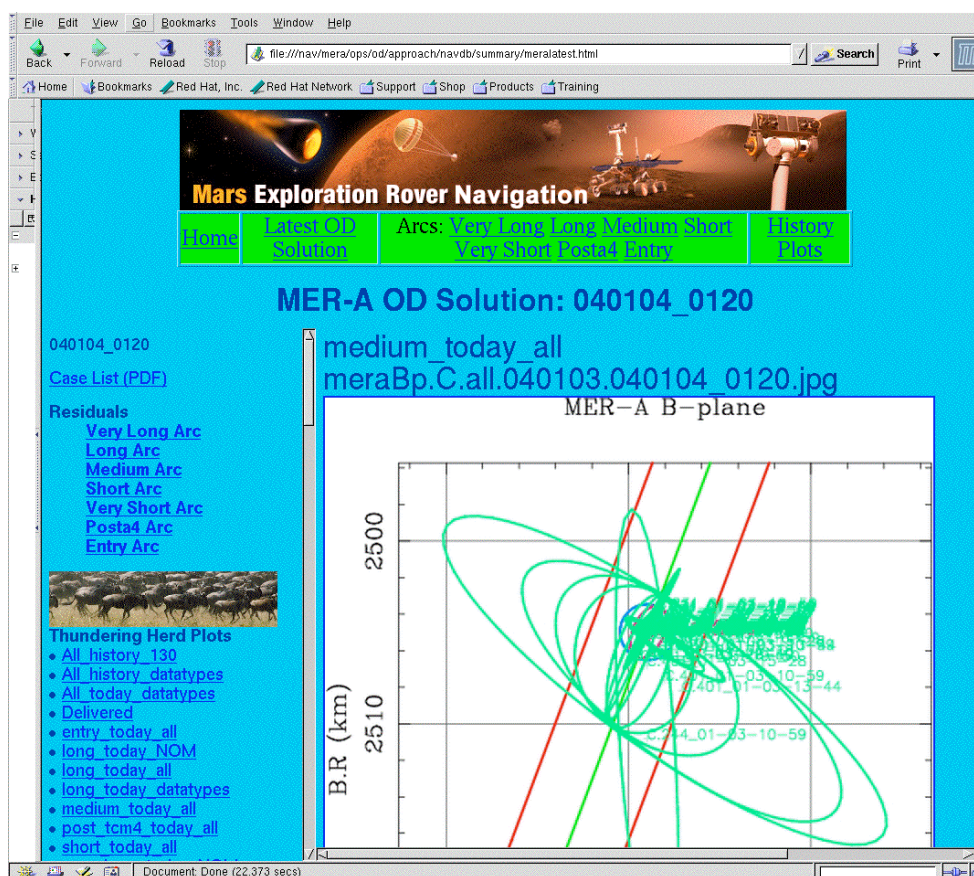


Figure 15. Typical "NavReport" web page.

The NAG consists of all section management staff in the relevant disciplines, as well as other interested and specifically invited parties, depending on the subject of the meeting.

deliveries, with the expectation that the deliveries before the next Daily Show would follow the recommendations from the meeting.

To make the Daily Show meetings efficient, an attempt was made to follow the agenda below (although this agenda was never explicitly displayed in the meetings). The meeting took place in a room with video projection capability, and so viewing NavReport web pages and xbpz B-plane plots was the starting point for many of the discussions.

Daily Show orbit determination agenda:

- News from last 24 hours – data outages, action item responses if any
- Data evaluation – last data cutoff, new passes and Δ DOR points, anything unusual in data, pre-fit and post-fit residuals from baseline case
- Case news – newly added cases, which cases were run, which arcs were run.
- B-plane plots – pre-generated plots on NavReport page, xbpz to look at plots interactively
- Navreport items – history plots, consistency plots, data statistics plots, anything else of interest
- Special topics if any
- Decide on any cases to add
- Decide on likely delivery case and arc for tomorrow

Following the orbit determination part of the meeting, maneuver and entry, descent, and landing issues were discussed. The Daily Show proved useful in focusing the Navigation team's thinking about orbit determination issues, and also served as a forum for MER project management to receive a greater understanding of the navigation process than could be provided in project-wide meetings.

V. Results

Since the overall navigation and orbit determination results are discussed in Refs. 1 and 2, the discussion here will be confined to specific items that are of interest to the filter strategy. This is particularly possible since there were no significant problems, both as a result of much preparation and because the proverbial Murphy was evidently busy elsewhere. The changes that were made to respond to events during the approach phase of both missions were greatly aided by the model and tool development, and in the absence of these preparations might have proved difficult.

One interesting general result is that during the last two months of the MER cruise trajectory, a total of 12185 orbit solutions were produced, using 108 different filter cases and 7 different data arcs for each vehicle. By comparison, the previous Mars approach for a NASA spacecraft (Mars Odyssey), while admittedly only consisting of one vehicle, performed less than a thousand orbit solutions during a similar time span. It will be interesting to see how this trend continues, and it may turn out that generating fewer solutions will be more appropriate for future spacecraft.

The detailed results described below have been alluded to already, and consist of a discussion of non-gravitational modeling accuracy and post-TCM-4 solution history, both for maneuver modeling variations and in general.

A. Non-gravitational Modeling Accuracy

One of the most comforting features of the MER approaches was the accuracy of the trajectory predictions. Figure 16 shows the delivered solutions for MER-B for the 5 weeks before TCM-B4 (with the solution data cutoff as the label). The December 12 solution was still using the looser attitude turn modeling without a non-zero axial component, but it still easily encompasses the later solutions. The remaining solutions are within half a kilometer, which suggests that the statistical error model of the earlier solutions is still too large.




Figure 16. B-plane plots of MER-B solutions, showing prediction consistency.

One of the biggest contributors to the non-gravitational modeling accuracy was the solar pressure model. Figure 17 shows the stochastic diffuse solar pressure estimates for the MER-B very-long arc. The process noise for this parameter is 0.033 (in JPL units, which is the same as 0.1 as a coefficient). The estimated values are never more than half a post-fit sigma from zero, although there is a clear offset of about 0.005 around the turn on November 20. As for the other significant non-gravitational model, none of the MER-B attitude turn ΔV estimates differed from the *a priori* value by more than 0.022 mm/sec, or 20 percent of the *a priori* sigma (which was 0.1 mm/sec, since the largest difference was in the axial direction).




Figure 17. Diffuse solar pressure stochastic estimates and sigmas for MER-B very-long arc.

B. Post TCM-A4 Solution History

Following TCM-A4, the evolution of B-plane estimates based on varying error models for the maneuver parameters was a subject studied with great interest, in an attempt to predict the ultimate maneuver delivery. Figure 18 shows solutions for data cutoffs between December 29, 18:30 and January 1, 13:00 (the data cutoff for TCM-A5). All of these solutions are statistically consistent with each other, and the original baseline and looser solutions moved toward the original tight solution.

Following the cancellation of TCM-A5, the B-plane estimates started to reflect sensitivity to the presence of Mars. Figure 19 shows a subset of the delivered solutions after TCM-A4 in white (before the TCM-A5 data cutoff) and red (after the TCM-A5 data cutoff). The progression of delivered solutions follows a rough semicircle over this time, and the B-plane uncertainty becomes highly constrained in the B-magnitude direction. The solutions




Figure 18. B-plane plot showing solution variations following TCM-A4.




Figure 19. B-plane plot showing delivered solutions and ephemeris estimation solutions after TCM-A4

estimating the ephemeris (shown in green) follow a roughly upward trend. All of the ephemeris estimates start from the original ephemeris *a priori* sigma, which is twice the size of the considered ephemeris uncertainties in the delivered solutions. From these results, it seems likely that an out-of-plane Mars ephemeris error of 400-500 meters

could have been present. The final delivered solution actually has a larger semi-major axis than the preceding one, due to consider effects (probably from the Mars ephemeris).

C. Post TCM-B4 Solution History

As discussed above, TCM-B4 was significantly larger and more complicated (both lateral and axial, with 2 lateral pulses) than TCM-A4. Figure 20 shows solutions for data cutoffs between January 18, 12:15 and January 22, 14:00 (the data cutoff for TCM-B5). Unlike the solutions following TCM-A4, the tight *a priori* sigmas for maneuver parameters did not work well, which is unsurprising in retrospect since some maneuver parameters changed by more than one sigma. The baseline, loose, and very loose solutions all eventually converge, and the solutions with the stochastic diffuse solar pressure constrained do the best job of predicting the ultimate B-plane location. Consequently, the constrained solar pressure solutions were delivered for several days following TCM-B4.

Figure 21 shows a subset of the solutions delivered after TCM-B4 in white (before the TCM-B5 data cutoff) and red (after the TCM-B5 data cutoff). The first two delivered solutions did not have the solar pressure constraint, and did not have enough data to determine the TCM-B4 maneuver parameters correctly. As with the MER-A approach, the solutions estimating the ephemeris move up in B•R, but the solutions considering the ephemeris move sharply downward at the end. In both types of solutions the ellipse semi-minor axis shrinks markedly as Mars approaches, but unlike MER-A approach, the semi-minor axis direction is not closely aligned with the B-magnitude direction, and consequently the entry flight path angle sigmas are larger for MER-B than for MER-A.




Figure 20. B-plane plot showing solution variations following TCM-B4.




Figure 21. B-plane plot showing delivered solutions and ephemeris estimation solutions after TCM-B4

VI. Conclusion

The development of the MER orbit determination strategy paid an unprecedented amount of attention to detail. This effort paid off in both accurate orbit determination and high confidence in the orbit determination process on the part of the project. The methods described here have become the standard for orbit determination strategy for future missions. Nonetheless, further improvements in the models and techniques are possible, and remain as a challenge for future navigation teams.

Appendix

1. A Statistical Measure of the Distance between two OD Solutions.


Navigation analysts have long used geometry to portray the uncertainty associated with an OD solution. Given a solution and covariance, we routinely map that solution and covariance to some point of interest such as the B -plane intercept and time to that intercept. The uncertainty of the mapping is represented graphically as the $1-\sigma$ ellipse (or ellipsoid). Mathematically, if \mathbf{c} is the solution, \mathbf{P} its covariance and $\mathbf{\Lambda}$ the information matrix ($\mathbf{\Lambda} = \mathbf{P}^{-1}$), then the quadratic “surface”

$$(\mathbf{x}-\mathbf{c})^T \mathbf{\Lambda} (\mathbf{x}-\mathbf{c}) = 1$$


is the $1-\sigma$ ellipsoid around the about the solution \mathbf{c} . More generally the $n-\sigma$ ellipsoid (where n is any positive real number) is the surface given by

$$(\mathbf{x}-\mathbf{c})^T \mathbf{\Lambda} (\mathbf{x}-\mathbf{c}) = n^2$$


Graphically, we often depict just the cross section of this ellipsoid with the B -plane.



When more than one solution is computed, the $1-\sigma$ ellipses are often plotted together. In such a situation, the question naturally arises “How close are these solutions to one another. This note presents a geometric technique for measuring the closeness between different solutions.



Suppose that $\mathbf{S}_1 = (\mathbf{c}_1, \mathbf{P}_1)$ and $\mathbf{S}_2 = (\mathbf{c}_2, \mathbf{P}_2)$ are two OD solutions. If $\mathbf{c}_1 \neq \mathbf{c}_2$ then for any positive t we can create $t-\sigma$ ellipsoids $\mathbf{E}_1(t)$ and $\mathbf{E}_2(t)$ about each solution respectively. For t sufficiently small, the surface and interiors of $\mathbf{E}_1(t)$ and $\mathbf{E}_2(t)$ do not overlap. As t increases from $\varepsilon > 0$ there will be a smallest t such $\mathbf{E}_1(t)$ and $\mathbf{E}_2(t)$ first overlap. At this value of t the two $t-\sigma$ are tangent. We call this first t the closeness measure of the two solutions and denote it by $d(\mathbf{S}_1, \mathbf{S}_2)$. Pictorially, we start with the $\varepsilon-\sigma$ ellipsoids and allow them to expand by linear scaling until they first touch. The point of tangency between the two ellipsoids has the same likelihood in both solutions. The number of σ of this tangency point from either solution is the measure of the closeness of the two solutions.



Mathematically, we want to find the smallest value of t such that

$$(\mathbf{x}-\mathbf{c}_1)^T \mathbf{\Lambda}_1 (\mathbf{x}-\mathbf{c}_1) = t^2 = (\mathbf{x}-\mathbf{c}_2)^T \mathbf{\Lambda}_2 (\mathbf{x}-\mathbf{c}_2) \quad (1)$$

has a exactly one solution \mathbf{x} . At the first point of tangency, the normals to the two ellipsoids will be anti-parallel. This is equivalent to the existence of a positive parameter α such that

$$\mathbf{\Lambda}_1 (\mathbf{x}-\mathbf{c}_1) = -\alpha \mathbf{\Lambda}_2 (\mathbf{x}-\mathbf{c}_2)$$

Solving for \mathbf{x} we have

$$\mathbf{x} = (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} (\mathbf{\Lambda}_1 \mathbf{c}_1 + \alpha \mathbf{\Lambda}_2 \mathbf{c}_2)$$

Substituting this value for \mathbf{x} back into equation (1), noting $\mathbf{\Lambda}_i$ is symmetric, and applying routine simplification yields:

$$\begin{aligned} \alpha^2 (\mathbf{c}_2-\mathbf{c}_1)^T \mathbf{\Lambda}_2 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_1 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_2 (\mathbf{c}_2-\mathbf{c}_1) = \\ (\mathbf{c}_2-\mathbf{c}_1)^T \mathbf{\Lambda}_1 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_2 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_1 (\mathbf{c}_2-\mathbf{c}_1) \end{aligned} \quad (2)$$

After making the observation that

$$(\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2) = \mathbf{\Lambda}_1 (\mathbf{P}_2 + \alpha \mathbf{P}_1) \mathbf{\Lambda}_2 = \mathbf{\Lambda}_2 (\mathbf{P}_2 + \alpha \mathbf{P}_1) \mathbf{\Lambda}_1$$

both sides of equation (2) can be re-written in terms of covariances:

$$\begin{aligned} \alpha^2 (\mathbf{c}_2-\mathbf{c}_1)^T \mathbf{\Lambda}_2 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_1 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_2 (\mathbf{c}_2-\mathbf{c}_1) \\ = \alpha^2 (\mathbf{c}_2-\mathbf{c}_1)^T (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} \mathbf{P}_1 (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} (\mathbf{c}_2-\mathbf{c}_1) \end{aligned}$$

and

$$\begin{aligned} (\mathbf{c}_2-\mathbf{c}_1)^T \mathbf{\Lambda}_1 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_2 (\mathbf{\Lambda}_1 + \alpha \mathbf{\Lambda}_2)^{-1} \mathbf{\Lambda}_1 (\mathbf{c}_2-\mathbf{c}_1) \\ = (\mathbf{c}_2-\mathbf{c}_1)^T (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} \mathbf{P}_2 (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} (\mathbf{c}_2-\mathbf{c}_1) \end{aligned}$$

Substituting these back into equation (2) and rearranging terms yields

$$0 = (\mathbf{c}_2 - \mathbf{c}_1)^T (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} (\alpha^2 \mathbf{P}_1 - \mathbf{P}_2) (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} (\mathbf{c}_2 - \mathbf{c}_1)$$

Let \mathbf{M}^C stand for the matrix of cofactors of a matrix \mathbf{M} and recall that for a symmetric matrix $\det(\mathbf{M}) \mathbf{M}^{-1} = \mathbf{M}^C$. Applying this observation gives

$$0 = (\mathbf{c}_2 - \mathbf{c}_1)^T (\mathbf{P}_2 + \alpha \mathbf{P}_1)^C (\alpha^2 \mathbf{P}_1 - \mathbf{P}_2) (\mathbf{P}_2 + \alpha \mathbf{P}_1)^C (\mathbf{c}_2 - \mathbf{c}_1) \quad (3)$$

Equation (3) is a polynomial equation in α of degree $2n$ where n is the dimension of \mathbf{c}_i . The positive real roots of this equation, $\{\alpha_0, \alpha_0, \dots, \alpha_k\}$ provide values of α for which the t - σ ellipsoids are tangent. Of these roots the one that minimizes

$$(\mathbf{c}_2 - \mathbf{c}_1)^T (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} \mathbf{P}_2 (\mathbf{P}_2 + \alpha \mathbf{P}_1)^{-1} (\mathbf{c}_2 - \mathbf{c}_1)$$

is the value associated with the first tangency condition. The square root of this value yields the measure of closeness of the two solutions.

$$d(\mathbf{S}_1, \mathbf{S}_2) = \sqrt{\min_{i=0, \dots, k} \left\{ (\mathbf{c}_2 - \mathbf{c}_1)^T (\mathbf{P}_2 + \alpha_i \mathbf{P}_1)^{-1} \mathbf{P}_2 (\mathbf{P}_2 + \alpha_i \mathbf{P}_1)^{-1} (\mathbf{c}_2 - \mathbf{c}_1) \right\}}$$

2. Properties of the compatibility measure

The closeness measure, d , has several numerical properties that lend support to the notion that it provides a useful measure of the distance between two OD solutions

1. $d(\mathbf{S}_1, \mathbf{S}_2) = d(\mathbf{S}_2, \mathbf{S}_1)$
2. For $t > 0$, $d((\mathbf{c}_1, \mathbf{P}_1), (\mathbf{c}_2, \mathbf{P}_2)) = d((t\mathbf{c}_1, t\mathbf{P}_1), (t\mathbf{c}_2, t\mathbf{P}_2))$
3. For $t > 0$, $d((t\mathbf{c}_1, \mathbf{P}_1), (t\mathbf{c}_2, \mathbf{P}_2)) = t d((\mathbf{c}_1, \mathbf{P}_1), (\mathbf{c}_2, \mathbf{P}_2))$
4. For $t > 0$, $d((\mathbf{c}_1, t\mathbf{P}_1), (\mathbf{c}_2, t\mathbf{P}_2)) = (1/t) d((\mathbf{c}_1, \mathbf{P}_1), (\mathbf{c}_2, \mathbf{P}_2))$
5. If \mathbf{P} and \mathbf{Q} are both positive definite $n \times n$ positive definite matrices and $\mathbf{v}^T \mathbf{P} \mathbf{v} < \mathbf{v}^T \mathbf{Q} \mathbf{v}$ for all non-zero \mathbf{v} , then $d((\mathbf{c}_1, \mathbf{P}), (\mathbf{c}_2, \mathbf{P}_2)) > d((\mathbf{c}_1, \mathbf{Q}), (\mathbf{c}_2, \mathbf{P}_2))$.

It is useful to interpret these mathematical statements.

1. The measure of compatibility does not depend upon the order of the two solutions; there is not a “preferred” solution.
2. The measure of compatibility is independent of scale. If we change units of the solution and covariance, we obtain the same measure.
3. If solutions are scaled, the measure scales accordingly.
4. If we scale the size of the covariance, the measure of the compatibility scales inversely. In other words, if knowledge of the solution decreases by a scale factor t , the compatibility of the solutions increases by the scale factor $1/t$.
5. If the uncertainty of a solution is increased, that solution is more compatible with other solutions.

Acknowledgments

The lead author, who also served as the MER Navigation team Deputy Team Chief for Orbit Determination, would like to gratefully acknowledge the many contributions of the rest of the MER Navigation team, and of others at JPL who provided products, services, and advice. The sacrifices of the families of the authors, both during the many inconvenient operationally-required absences, and during the writing of this and other papers, is also much appreciated.

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

References

¹ D'Amario, L. A., "Mars Exploration Rovers Navigation Results," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-4980, AIAA, Washington, DC, 2004.

² Portock, B. M., Graat, E. J., McElrath, T. P., Watkins, M. M., and Wawrzyniak, G. G., "Mars Exploration Rovers Cruise Orbit Determination," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-4981, AIAA, Washington, DC, 2004.

³ Ratcliff, J. T., "KEOF Operational EOP Deliveries During MER," 24 July 2003 (JPL internal memo).

⁴ Wawrzyniak, G. G., Baird, D. T., Graat, E. J., McElrath, T. P., Portock, B. M., and Watkins, M. M., "Mars Exploration Rovers Orbit Determination System Modeling," *AIAA/AAS Astrodynamics Specialist Conference*, Providence, AIAA-2004-4983, AIAA, Washington, DC, 2004.

⁵ Potts, C. L., Raofi, B., and Kangas, J. A., "Mars Exploration Rovers Propulsive Maneuver Design," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-4985, AIAA, Washington, DC, 2004.